

The Neutral Hydrogen Content of Galaxies in Cosmological Hydrodynamic Simulations

Romeel Davé^{1,2}, Neal Katz³, Benjamin D. Oppenheimer⁴, Juna A. Kollmeier⁵, David H. Weinberg⁶

¹ University of the Western Cape, 7535 Bellville, Cape Town, South Africa

² Astronomy Department, University of Arizona, Tucson, AZ 85721, USA

³ Astronomy Department, University of Massachusetts, Amherst, MA 01003, USA

⁴ Leiden Observatory, Leiden University, PO Box 9513, 2300 RA Leiden, Netherlands

⁵ Observatories of the Carnegie Institute of Washington, Pasadena, CA 91101, USA

⁶ Astronomy Department and CCAPP, Ohio State University, Columbus, OH 43210, USA

18 February 2013

ABSTRACT

We examine the global H_I properties of galaxies in quarter-billion particle cosmological hydrodynamic simulations, focusing on how our main adjustable physical process, galactic outflows, impacts H_I content. In addition to the three outflow models considered in our earlier papers, we present a new one (ezw) motivated by high resolution interstellar medium simulations, in which the scalings of wind speeds and mass loading factors follow those expected for momentum-driven outflows for larger galaxies, and energy-driven outflows for dwarf galaxies ($\sigma < 75 \text{ km s}^{-1}$). To obtain predicted H_I masses, we employ a simple but effective local correction for particle self-shielding, as well as an observationally-constrained transition from neutral to molecular hydrogen. We find that our ezw model produces an H_I mass function whose shape agrees well with observations from the ALFALFA survey, having a low mass end slope of -1.3 , while other models agree less well. Outflows critically govern the H_I content in low-mass galaxies, with higher mass loading factors yielding higher H_I fractions. Satellite galaxies have a bimodal distribution in H_I fraction versus halo mass, with lower mass satellites and/or satellites in larger halos more often being devoid of H_I. At a given stellar mass, H_I content correlates with star formation rate and inversely correlates with metallicity, as expected if driven by stochasticity in the accretion rate. At higher redshifts, massive H_I galaxies become less frequent and the H_I mass function becomes significantly steeper. The global cosmic H_I density conspires to remain fairly constant from $z \sim 5 \rightarrow 0$, but the relative contribution from smaller galaxies increases substantially with redshift. Overall, H_I in galaxies reflects a transient reservoir of fuel for star formation, and hence provides a crucial glimpse into the inflow and outflow processes that govern galaxy evolution.

1 INTRODUCTION

The current paradigm for galaxy evolution rests on the tenet that gas flows into and out of galaxies are primarily responsible for governing galaxy growth (e.g. Kereš et al. 2005; Dekel et al. 2009; Davé, Finlator, & Oppenheimer 2012, and references therein). Such accretion and outflow processes are often difficult to detect directly owing to the tenuous and multi-phase nature of the gas, but they are expected to leave clear imprints on the stellar and gaseous content of galaxies. A key region where such baryon cycling processes might be probed is in the outskirts of star-forming galaxies, where reservoirs of neutral hydrogen (H_I) are expected to hold the raw material that ultimately fuels star formation. This reservoir is thought to be continually replenished by gravitationally-driven accretion from the highly

ionised intergalactic medium (IGM), possibly augmented by re-accretion of ejected gas.

Observations of the neutral hydrogen content of galaxies using 21cm fine structure line emission are progressing rapidly, and promise to improve further. The H_I Parkes All-Sky Survey (HIPASS; Meyer et al. 2004) provided a comprehensive deep census of the H_I content of low-redshift galaxies. This was followed by the Arecibo Fast Legacy Alfa Survey (ALFALFA; Giovanelli et al. 2005), which surveyed 7000 deg² for H_I 21cm emitting objects. More recently, the GALEX Arecibo SDSS Survey (GASS; Catinella et al. 2010) has obtained high-fidelity H_I data on optically-selected galaxies, to provide a more H_I-unbiased census. Construction is underway of the MeerKAT telescope, which will provide the ability to detect H_I in galaxies out to unprecedented redshifts (e.g. the Looking At the Distant Universe with the MeerKAT Array – LADUMA – Survey; Holwerda et al.

2011), and the Australian SKA Pathfinder (ASKAP), which will analogously perform the Wallaby Survey. Ultimately, the Square Kilometer Array (SKA) will enable studies of H_I both nearby and out to higher redshifts that far exceed today’s capabilities. Complementary information is provided by studies of damped Ly α absorbers (DLAs) that trace H_I in and around galaxies in the spectra of background quasars (e.g. Battisti et al. 2012).

It is therefore timely and crucial to develop a theoretical framework for understanding the physics that governs the H_I content of galaxies and its evolution. Because galactic H_I reservoirs are thought to represent a transient phase of baryons (Prochaska & Wolfe 2009) as they pass from the diffuse intergalactic medium (e.g. Davé et al. 2010) to higher density molecular gas, which eventually converts into stars, it is essential that models for cosmic H_I be dynamical in nature. Cosmological hydrodynamic simulations provide one such type of dynamical model for this purpose, allowing one to directly track the physical state of gas as it flows through the H_I reservoir. Particularly, the inclusion of galactic outflows in simulations provides some interesting new twists on the accretion paradigm, as enriched outflows that return back into galaxies may provide a significant source of accretion (“wind recycling”; Oppenheimer et al. 2010). These inflow, outflow, and recycling processes are expected to manifest themselves in the H_I content of galaxies.

In this work, we examine the H_I content of galaxies in cosmological hydrodynamic simulations with a variety of outflow models. This work follows on our early study in Popping et al. (2009), but it uses a significantly improved simulation suite and focuses on the physical processes that govern the H_I content in addition to making predictions for H_I observables and their evolution. This work is also comparable to that of Duffy et al. (2012), who used simulations to examine both the atomic and molecular content of simulated galaxies. Compared to that work, our model explores a different range of feedback parameters (including ones that match H_I observations substantially better), probes down to significantly smaller systems (albeit in a smaller volume), and uses a slightly different approach to compute the H_I content of galaxies. Nonetheless, for overlapping models and mass ranges, our results generally agree with those of Duffy et al. (2012).

The paper is organised as follows. In §2 we introduce our simulations, and describe our method for calculating the atomic and molecular gas content of our simulated galaxies. In §3 we examine the H_I mass function and its evolution, and we examine the complementary constraint of H_I richness versus stellar mass in §4. In §5 we study how H_I content is impacted by environment, and in §6 we show how it relates to galaxy metallicity and star formation rate. In §7 we make predictions for the evolution of Ω_{HI} , and compare to observations out to $z \sim 4$. In §8 we examine key H_I properties in a simulation with different numerical resolution to test for convergence. Finally, we summarise and discuss our results in §9.

2 METHODS

2.1 Simulations

Our simulations are evolved with an extended version of the GADGET-2 N-body + Smoothed Particle Hydrodynamic (SPH) code (Springel 2005). We assume a Λ CDM cosmology (Hinshaw et al. 2009): $\Omega_M = 0.28$, $\Omega_\Lambda = 0.72$, $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1}) = 0.7$, a primordial power spectrum index $n = 0.96$, an amplitude of the mass fluctuations scaled to $\sigma_8 = 0.82$, and $\Omega_b = 0.046$. We call this cosmology our r-series, where our general naming convention is $r[\text{boxsize}]n[\text{particles}/\text{side}][\text{wind model}]$. Our primary simulations use a cubic volume of $32h^{-1}\text{Mpc}$ on a side with 512^3 dark matter and 512^3 gas particles, and a softening length of $\epsilon = 1.25h^{-1}\text{kpc}$ (comoving, Plummer equivalent). The gas particle mass is $4.5 \times 10^6 M_\odot$, and the dark matter particles masses are approximately five times larger. We can thus reliably resolve galaxies down to stellar masses of $M_{*,\text{lim}} = 1.4 \times 10^8 M_\odot$ (see discussion below).

Our version of GADGET-2 includes cooling processes using the primordial abundances as described by Katz et al. (1996), with additional cooling from metal lines assuming photo-ionisation equilibrium from Wiersma et al. (2009). Star formation is modelled using a subgrid recipe introduced by Springel & Hernquist (2003) where a gas particle above a density threshold of $n_H = 0.13 \text{ cm}^{-3}$ is modelled as a fraction of cold clouds embedded in a warm ionised medium following McKee & Ostriker (1977). Star formation (SF) follows a Schmidt law (Schmidt 1959) with the SF timescale scaled to match the $z = 0$ Kennicutt relation (Kennicutt 1998). We use a Chabrier (2003) initial mass function (IMF) throughout. We account for metal enrichment from Type II supernovae (SNe), Type Ia SNe, and AGB stars, and we track four elements (C,O,Si,Fe) individually, as described by Oppenheimer & Davé (2008).

Galactic outflows are implemented using a Monte Carlo approach analogous to star formation. Outflows are directly tied to the SFR, using the relation $\dot{M}_{\text{wind}} = \eta \times \text{SFR}$, where η is the outflow mass loading factor. The probability for a gas particle to spawn a star particle is calculated from the subgrid model described above, and the probability to be launched in a wind is η times the star formation probability. If the particle is selected to be launched, it is given a velocity boost of v_w in the direction of $\mathbf{v} \times \mathbf{a}$, where \mathbf{v} and \mathbf{a} are the particle’s instantaneous velocity and acceleration, respectively. Once a gas particle is launched, its hydrodynamic (not gravitational) forces are turned off until either $1.95 \times 10^{10}/(v_w \text{ (km s}^{-1}))$ years have passed or, as more often occurs, the gas particle has reached a density that is 10% of the SF critical density (i.e., 0.013 cm^{-3}). This attempts to mock up chimneys generated by outflows that would allow a relatively unfettered escape from the galactic ISM, a process not properly captured by the spherically-averaging SPH algorithm at $\sim\text{kpc}$ resolution. It also yields results that are less sensitive to numerical resolution (Springel & Hernquist 2003b) than models that do not turn off hydrodynamic forces. For a further discussion of hydrodynamic decoupling, see Dalla Vecchia & Schaye (2008).

Choices of the parameters η and v_w define the “wind model”. For this paper we make use of the following four wind models:

(i) **No winds (nw)**, where we do not include outflows

(i.e. $\eta = 0$). This model fails to match a wide range of observables (e.g. Davé, Oppenheimer, & Finlator 2011; Davé, Finlator, & Oppenheimer 2011), but is included to establish a baseline for the overall impact of winds.

(ii) **Constant winds (cw)**, where $\eta = 2$ and $v_w = 680 \text{ km s}^{-1}$ for all galaxies. This model is similar to the wind model employed by Duffy et al. (2012).

(iii) **Momentum-conserving winds (vzw)**, where the wind speed and the mass loading factor depend on the galaxy velocity dispersion σ (Murray, Quataert, & Thompson 2005), scaling as estimated (see Oppenheimer & Davé 2008).

$$v_w = 3\sigma\sqrt{f_L - 1}, \quad (1)$$

$$\eta = \frac{\sigma_0}{\sigma}, \quad (2)$$

where $f_L = [1.05, 2]$ is the luminosity factor in units of the galactic Eddington luminosity (i.e. the critical luminosity necessary to expel gas from the galaxy potential), and $\sigma_0 = 150 \text{ km/s}$ is the normalisation of the mass loading factor. Choices for the former are taken from observations (Rupke, Veilleux & Sanders 2005), while the latter is broadly constrained to match high-redshift IGM enrichment (Oppenheimer & Davé 2008). Galaxies are identified “on-the-fly” during the simulation using a Friends-of-Friends (FOF) algorithm applied to the gas, star, and dark matter particles. We choose a smaller than usual linking length to pick out just the galaxies. This linking length evolves with redshift and in terms of the mean interparticle separation of all particles is

$$0.06\left(\frac{H(z)}{H_0}\right)^{1/3}. \quad (3)$$

We estimate the velocity dispersions necessary for our wind scaling laws using the total FOF galaxy mass M_{gal}^1 using the relation

$$\sigma = 200\left(\frac{M_{\text{gal}}}{5 \times 10^{12} h^{-1} M_\odot}\right)^{1/3} \text{ km/s}, \quad (4)$$

Which we have empirically determined gives an accurate measure of the velocity dispersion in our simulations.

(iv) **Hybrid energy/momentunm-driven winds (ezw)**, which employs the vzw scalings for galaxies with $\sigma > 75 \text{ km s}^{-1}$, then switches over to a steeper dependence of $\eta \propto \sigma^{-2}$ in $\sigma < 75 \text{ km s}^{-1}$ systems. The wind speed still scales proportionally to σ as in the vzw model. This model roughly captures the behaviour in recent analytic and hydrodynamic models of outflows from interstellar media by Murray, Quataert, & Thompson (2010) and Hopkins, Quataert, & Murray (2012), respectively. The basic idea is that in dwarf galaxies the energy from supernovae plays a dominant role in driving outflows, while in larger systems the momentum flux from young stars and/or supernovae is the dominant driver. As a result, the outflow scalings switch from momentum-driven at high masses to energy-driven at low masses. We make this transition abruptly at $\sigma = 75 \text{ km s}^{-1}$ guided by the analytic models

¹ We have actually been using this algorithm, which slightly differs from (Oppenheimer & Davé 2008), in all of our recent simulations starting with (Oppenheimer et al. 2010).

of Murray, Quataert, & Thompson (2005, 2010) and high-resolution ISM simulations of Hopkins, Quataert, & Murray (2012), although it should perhaps be more gradual (note that η itself is continuous across this boundary). In any case, this model captures the gist of the most up-to-date small-scale outflow models. As we will see, the ezw outflow model fares somewhat better compared to observations than vzw, which in turn fares much better than cw or nw.

Additionally, in our fiducial ezw simulation, we employ a heuristic prescription to quench star formation in massive galaxies. This is not a physical model, but simply a tuned parametrisation to limit star formation in massive systems to reproduce the observed exponential high-mass cutoff in the stellar mass function and to make the simulation faster to run. In this prescription, we quench star formation in an *entire galaxy* according to a probability, P_Q , given by the equation

$$P_Q = 1 - \frac{1}{2} \text{erfc}\left[\frac{\log(\sigma) - \log(\sigma_{Q\text{med}})}{\log(\sigma_{Q\text{spr}})}\right], \quad (5)$$

where the median σ at which a galaxy has a 50% chance of being quenched is $\sigma_{Q\text{med}} = 110 \text{ km s}^{-1}$, corresponding to $M_{\text{halo}} = 10^{12.1} M_\odot$ at $z = 0$, and the spread in σ is $\sigma_{Q\text{spr}} = 32 \text{ km s}^{-1}$. We also require $\sigma > 75 \text{ km s}^{-1}$ to fully suppress the already low probability of lower mass galaxies being quenched. We have found that these parameter choices nicely reproduces the high-mass end of the stellar mass function, as we will show in §2.3.

Every time we identify galaxies using our FOF group finder, which we do to calculate σ for the vzw and ezw wind models and occurs approximately every 10 Myr, we compute this quenching probability for each identified galaxy. If the galaxy is “quenched”, then each time a gas particle would have formed a star over the time interval until the next time we identify galaxies using our FOF group finder, it is instead heated to 50 times the virial temperature, T_{vir} , where

$$T_{\text{vir}} = 5 \times 10^6 \left(\frac{\sigma}{200 \text{ km s}^{-1}}\right)^2 \text{ K}. \quad (6)$$

The motivation for heating the gas to such extreme temperatures is primarily to prevent it from re-accreting at later times and thus requiring multiple ejections; the total energy input is thus less for higher quenching temperatures. The median temperature to which the ISM is heated by quenching is $10^{8.1} \text{ K}$ and arises from a median halo mass of $10^{12.2} M_\odot$. The energy input from quenching averages to $10^{40.5} \text{ ergs}^{-1} \text{Mpc}^{-3}$ between $z = 0.75 - 2.5$, corresponding to the peak of AGN activity, and the integrated energy input until $z = 0$ is $8 \times 10^{14} \text{ ergs/g}$, which equals 9×10^{-7} of the rest mass energy of all baryons. Considering that 6.5% of baryons are in stars in this simulation at $z = 0$, and using the assumption that supermassive black holes (SMBHs) have 10^{-3} of the mass in stars, the quenching energy corresponds to 1.4% of the rest-mass energy of SMBHs, which is comparable to the energy imparted from AGN feedback in cosmological simulations that self-consistently track black hole growth and feedback (Di Matteo et al. 2005; Booth & Schaye 2009).

In this paper we are mostly concerned with galaxies around L^* and below, which are mostly unquenched. We emphasise that this quenching prescription is not intended to be a realistic physical model, and is actually in large part a computational convenience, as removing gas from the largest galaxies substantially speeds up the code at low redshifts

and allows us to run our large simulations to $z = 0$ within practical time frames. We refer the reader to Duffy et al. (2012) for a discussion of H_I in larger-volume simulations that attempt to quench massive galaxies based on directly tracking black hole growth and feedback.

Finally, our no-wind simulation employs a volume of $16h^{-1}\text{Mpc}$ with 2×256^3 particles. This results in the same mass and spatial resolution as the wind simulations, but in a volume that is 8 times smaller. As we mostly use the no-wind simulation to qualitatively show the impact of winds at small masses, the reduced volume for simulation will not significantly affect our conclusions.

2.2 Computing the H_I content

We use SKID² (Spline Kernel Interpolative Denmax) to identify galaxies as bound groups of star-forming gas and stars (Kereš et al. 2005; Oppenheimer et al. 2010). Our galaxy stellar mass limit is set to be ≥ 64 star particles (Finlator et al. 2006), resulting in a minimum resolved mass of $M_{*,\min} = 1.4 \times 10^8 M_\odot$. We will only consider galaxies with stellar masses $M_* \geq M_{*,\min}$ in our analysis, regardless of their H_I content. Our resolution convergence tests in §8 show that the H_I properties are reasonably well converged even at this stellar mass threshold.

We identify dark matter halos via a spherical overdensity algorithm (Kereš et al. 2005) out to a virial overdensity given in Davé et al. (2010, see their eq. 1), centred on each galaxy. We separate galaxies into central and satellite galaxies by associating each galaxy with a halo; if a galaxy’s centre lies within the virial radius of a larger (by stellar mass) galaxy, we consider it to be a satellite of that galaxy, and the halos of those two galaxies are merged. Note that for galaxies near the edge of a larger halo, this can result in halo that has “bumps” along its (mostly spherical) surface.

To compute the H_I content of galaxies, we need to isolate the gas that is in the (predominantly) neutral phase. This involves defining two “boundaries”: the division between gas that is exposed to the full metagalactic ionising flux and gas that is self-shielded, and the division between atomic and molecular gas.

To calculate the self-shielding, we first compute the neutral hydrogen component of each gas particle under the assumption that it is *not* self-shielded. We follow Popping et al. (2009) who employed a simple hydrogen ionisation balance calculation that yielded the following formula to determine the neutral fraction:

$$f_{\text{HI}} = \frac{2C + 1 - \sqrt{(2C + 1)^2 - 4C^2}}{2C} \quad (7)$$

with

$$C = \frac{n\beta(T)}{\Gamma_{\text{HI}}} \quad (8)$$

where n is the Hydrogen number density, T is the gas temperature, Γ_{HI} is the H_I photo-ionisation rate, and the recombination rate coefficient β is well-fit by the func-

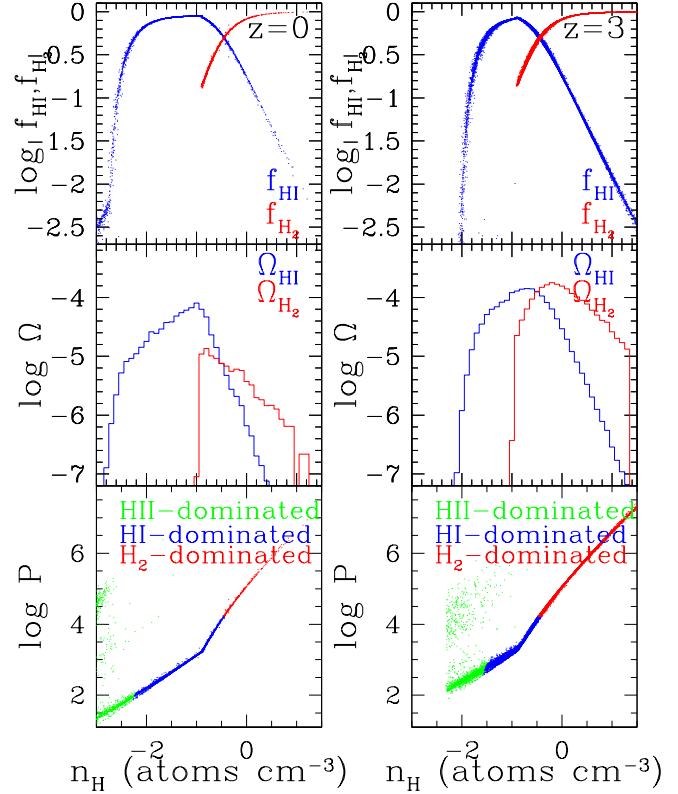


Figure 1. Gas particle properties as a function of number density n_H in our r48n384vzw simulation. Left panels show $z = 0$, right panels show $z = 3$. *Top:* H_I fraction (blue) and H₂ fraction (red) based on the calculation described in the text. We show a random sampling of 1% of all the particles. *Middle:* Total mass density Ω per unit $\log n_H$ in H_I (blue) and H₂ (red). *Bottom:* Pressure-density relation. Green points show unshielded particles (either below the pressure threshold or above the density threshold), blue points show H_I-dominated gas, and red points show molecular-dominated gas.

tion (Verner & Ferland 1996)

$$\beta(T) = a \left[\sqrt{(T/T_0)}(1 + \sqrt{(T/T_0)})^{1-b} (1 + \sqrt{(T/T_1)})^{1+b} \right]^{-1}. \quad (9)$$

For H_I, the fitting parameters are $a = 7.982 \times 10^{-11} \text{ cm}^3 \text{s}^{-1}$, $b = 0.7480$, $T_0 = 3.148 \text{ K}$, and $T_1 = 7.036 \times 10^5 \text{ K}$. We take Γ_{HI} from the Haardt & Madau (2001) ionising background, whose amplitude is adjusted to match the observed mean flux decrement in the Ly α forest in these simulations (see Davé et al. 2010); we describe this further below.

With the optically-thin neutral fraction computed for each gas particle, we employ a simple particle-by-particle post-processing correction for self-shielding of the H_I. We resort to this approach because it is computationally prohibitive to do the full radiative line transfer on these simulations; instead, we will calibrate our approach to radiative transfer simulations by Faucher-Giguere, Kereš, & Ma (2010).

We begin by assuming that each particle has a density

² <http://www-hpcc.astro.washington.edu/tools/skid.html>

profile given by the SPH spline kernel $W(r)$ (see Springel 2005, for definition). We compute the radial column density profile as follows:

$$N_{\text{HI}}(r) = \frac{0.76 f_{\text{HI}} \rho_g}{m_p} \int_r^h W(r') dr', \quad (10)$$

where ρ_g is the SPH density of the gas particle, m_p is the proton mass, and h is the particle's smoothing length. We then determine the radius R where $N_{\text{HI}}(R) = N_{\text{HI,lim}}$, with $N_{\text{HI,lim}} = 10^{17.2} \text{ cm}^{-2}$ being where the particle becomes optically thick ($\tau = 1$) to continuum photons at the Lyman limit. If no such radius R exists, then the particle is optically thin. If R exists, then we compute the unshielded and shielded mass fraction using $4\pi \int W(r) r^2 dr$ integrated from $h \rightarrow R$ and $R \rightarrow 0$, respectively. The neutral fraction of the particle is then the mass-weighted mean neutral fraction, assuming that the portion from $R \rightarrow 0$ is 90% neutral (since some ionised gas is seen to exist even within dense regions). By this procedure, particles in dense regions "auto-shield" themselves from the ambient flux, though the effects of shielding from nearby gas are neglected.

This computation relies on knowing the metagalactic photo-ionising flux, since this determines the neutral fraction in the optically-thin regime. Our simulations assume a Haardt & Madau (2001) ionising background, but more detailed constraints can be placed on the amplitude of the HI photo-ionisation rate by using the measured mean flux decrement in the Lyman alpha forest. In particular, at each redshift for each simulation, we determine a "flux factor", which is the value by which we must multiply the strength of the Haardt & Madau (2001) background to achieve consistency with the observed flux decrement. For the observed flux decrement, we employ at $z < 2$ the determination by Kirkman et al. (2007), while for $z \geq 2$ we take the values from Becker et al. (2012). As described more fully in Davé et al. (2010), we extract 1000 spectra and iteratively adjust the flux factor until the mean flux is within 1% of the observed value. The resulting flux factors for our ezw simulation at $z = 0, 1, 2, 3, 4, 5$ are 1.72, 2.03, 1.66, 0.92, 1.13, and 1.70, respectively. For our vzw and cw simulations, the $z = 0$ flux factors are 1.63 and 1.41 respectively (we do not consider these models at higher redshifts).

Next, we separate the molecular component from the atomic HI. To do so, we employ the observed ISM pressure relation from The HI Nearby Galaxy Survey (THINGS; Leroy et al. 2008); specifically, we use their "combined spiral subsample" fit, which gives the ratio of molecular to atomic gas as

$$R_{\text{mol}} = (P/P_0)^\alpha, \quad (11)$$

where $P_0 = 1.7 \times 10^4 \text{ cm}^{-3} \text{ K}$ and $\alpha = 0.8$. We compute the gas pressure based on the density and two-phase medium temperature of each star-forming gas particle (our prescription follows Springel & Hernquist 2003). Note that we only compute molecular fractions for star-forming gas particles, which in our simulations is gas with $n_H > 0.13 \text{ cm}^{-3}$. Gas that is not star-forming is assumed to have zero molecular content, and in any case the formula above would yield a very small molecular fraction. This prescription follows that employed by Duffy et al. (2012), but differs from Popping et al. (2009) who used a fixed pressure threshold of $810 \text{ cm}^{-3} \text{ K}$ to separate atomic from molecular gas.

Figure 1 illustrates the resulting gas particle atomic and molecular fractions. The top left panel shows the HI and H₂ mass fractions in particles from our momentum-driven wind scalings (vzw) simulations at $z = 0$. In the optically thin regime (mostly not depicted here; see Davé et al. 2010, for a full phase space diagram), f_{HI} scales linearly with n_H . Then there is a sharp upturn at $n_H \sim 2 \times 10^{-3} \text{ cm}^{-3}$, when auto-shielding becomes important.

At high redshifts, the stronger ionising background causes auto-shielding to set in at a higher density. A simple scaling shows that $N_{\text{HI,lim}} \propto f_{\text{HI}} n_H l$, where the length $l \propto n_H^{-1/3}$ for a spherical cloud (or SPH particle), and $f_{\text{HI}} \propto n_H / \Gamma_{\text{HI}}$; hence $N_{\text{HI,lim}} \propto n_H^{5/3} / \Gamma_{\text{HI}}$. Therefore, a factor of 10 increase in Γ_{HI} , which occurs between $z = 0 \rightarrow 3$, will lead to a factor of 4 increase in the n_H where self-shielding becomes important, because $n_H \propto \Gamma_{\text{HI}}^{3/5}$ for constant $N_{\text{HI,lim}}$. Figure 1 (top right panel) shows that auto-shielding at $z = 3$ is effective above $n_H \sim 10^{-2} \text{ cm}^{-3}$ as expected for the stronger $z = 3$ background UV field, which also matches the $z = 3$ radiative transfer simulations of Faucher-Giguere, Kereš, & Ma (2010, see their Figure 3). Indeed the overall shape of $f_{\text{HI}}(n_H)$ is actually quite similar to theirs, although there is less scatter at a given n_H in our prescription owing to the fact that we do not consider shielding from neighbouring particles. This shows that our physically-motivated choice of auto-shielding with $N_{\text{HI,lim}}$ at the Lyman limit is a reasonable approximation to much more detailed radiative transfer models.

Moving to higher densities, eventually we reach the star-forming density threshold, above which the molecular fraction becomes nonzero. At $n_H \gtrsim 0.5 \text{ cm}^{-3}$, corresponding to a pressure close to P_0 , the gas becomes molecular-dominated. The pressure relation is shown in the bottom panels of Figure 1, with the particles colour-coded by their dominant phase of hydrogen. The THINGS pressure threshold is much higher than the $P/k = 810 \text{ cm}^{-3} \text{ K}$ assumed in Popping et al. (2009) and, moreover, our prescription produces a more gradual transition between atomic and molecular-dominated gas. A change in slope at $n_H = 0.13 \text{ cm}^{-3}$ occurs because this is the density above which we allow star to form and the gas particles become two-phase.

In this paper we are mostly interested in the regime $10^{-2.5} \lesssim n_H \lesssim 0.5 \text{ cm}^{-3}$ where the majority of cosmic HI resides (at $z = 0$), above which gas becomes mostly molecular and below which it is mostly ionised. This is shown in the middle panels of Figure 1, where we plot the total mass density Ω in HI (blue) and H₂ (red) per unit interval of $\log n_H$; the peak is around $n_H \sim 10^{-1} \text{ cm}^{-3}$. This HI-dominant density regime is reasonably well resolved in our simulations, and hence the predicted HI content is expected to be robust, despite it being a transitory phase. In contrast, the transition from molecular gas to stars typically occurs at densities well above what we can resolve directly, and hence the molecular content will not be as robustly predicted; for this reason comparisons to the molecular content of galaxies may be less robust (though still interesting; we refer the reader to Duffy et al. 2012, for such predictions). Nonetheless, our predictions for the evolution from $z \sim 2 \rightarrow 0$ of the total star-forming gas (most of which is typically molecular) in our momentum-driven wind (vzw) simulation are in good agreement with observations from

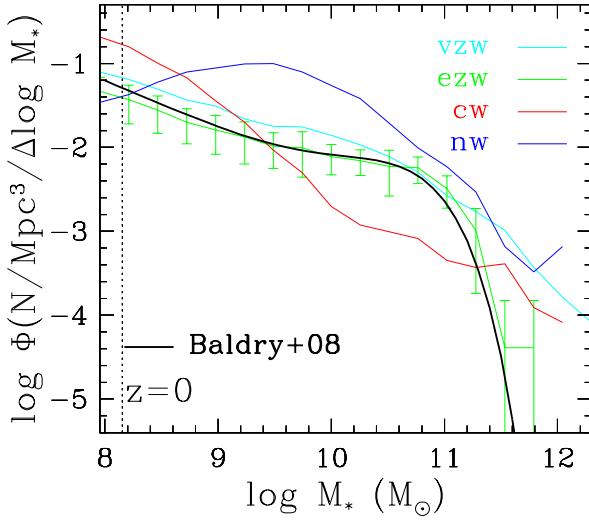


Figure 2. Galaxy stellar mass functions at $z = 0$ in our four wind models: ezw (green), vzw (cyan), cw (red), and nw (blue). The vertical dotted line is the stellar mass resolution limit. The thick solid line shows a fit to the observed GSMF from Baldry, Glazebrook, & Driver (2008). The ezw outflow model, which includes a heuristic prescription for quenching star formation at high masses, matches the observed GSMF to within uncertainties.

the Plateau de Bure High- z Blue Sequence Survey (PHIBSS; Tacconi et al. 2012, see their Figure 13); other wind models fare less well. Finally, we note that the stellar mass growth rate is fairly robust because it is not a transitory phase but an end state of accreted gas, and is typically limited by the gas supply rate (e.g. Finlator & Davé 2008; Bouché et al. 2010; Davé, Finlator, & Oppenheimer 2012).

With each gas particle’s H_I content determined, we must now associate the gas particles with galaxies. The information from SKID is not sufficient, because SKID only includes star-forming gas and stars in a galaxy, while a significant amount of the H_I resides in an extended region around the galaxy, beyond the actively star-forming region. Thus to account for extended H_I, we add up all the H_I mass in a sphere around each galaxy that extends to the outermost radius as defined by SKID, i.e. the radius of the farthest SKID particle associated with that galaxy. The outermost radius is typically many times the half-mass radius. While it may seem that this still may not fully account for an extended H_I disk, in practise the low threshold density for star formation in our simulations means that this choice still encompasses the vast majority of the H_I. We tried extending this to 1.5 times the outermost radius, and the total H_I mass in the volume increased by only 4%; furthermore, one will start to increasingly “double count” gas that may be between nearby galaxies as being part of both galaxies.

2.3 The Stellar Mass Function

As an initial baseline statistic to compare our four wind models, we show in Figure 2 their galaxy stellar mass functions (GSMFs) at $z = 0$. We compare these to observations from Baldry, Glazebrook, & Driver (2008) using SDSS.

The ezw model provides a strikingly good fit to the observed GSMF. At the massive end, this is a direct consequence of tuning the quenching prescription as described previously. In contrast, the low-mass end is unaffected by our quenching prescription, and instead reflects the effect of the hybrid energy-momentum driven winds. There is a clear upturn in the GSMF at $M_* \lesssim 10^{9.5} M_\odot$, which is also seen in the data. As explained in Oppenheimer et al. (2010), in our simulations this arises because above this mass, wind recycling becomes increasingly important, and provides extra fuel to higher mass galaxies. Below this mass, the typical recycling time becomes longer than a Hubble time, and the slope begins to steepen towards the dark matter halo mass function’s slope.

The vzw model provides not quite as good a fit to the low mass end, similar to what was seen in a lower-resolution version of the same model in Davé, Oppenheimer, & Finlator (2011). The differences at the massive end effectively show the impact of the quenching model, since our ezw simulation uses it while our vzw simulation does not, while the two outflow prescriptions themselves are identical in this mass regime. Quenching has a substantial effect on the GSMF, but we will show that it has a minimal effect on the H_I mass function. Meanwhile, the constant wind model produces a very steep low mass end slope that looks nothing like the data, while the no wind model overproduces stars at virtually all masses because it strongly overcools baryons. These results again follow those obtained using lower-resolution simulations of the same wind models presented in Oppenheimer et al. (2010) and Davé, Oppenheimer, & Finlator (2011), and more discussions of the GSMF can be found there.

This ezw simulation is the first hydrodynamic simulation that we are aware of to yield agreement with the observed GSMF to within statistical uncertainties over the entire mass range probed. This does not imply that this model is fully correct, as it could still be that the star formation histories in this model are incorrect, and other constraints may not be as well matched. As a case in point, we will show in §6 that the mass-metallicity relation in this model looks too steep compared to observations, and that the vzw model provides a better match. Furthermore, the success of this model at high masses owes to including an ad hoc and highly tuned prescription for quenching star formation that is not a direct implementation of a physical quenching mechanism. We leave a full comparison of the ezw model to a wide range of observables for future work, and focus here on the H_I properties of galaxies in our four wind simulations.

3 THE HI MASS FUNCTION

The most basic counting statistic for characterising the H_I content of galaxies is the H_I mass function (HIMF). Improving 21cm observations have enabled the HIMF to be probed down to H_I masses approaching $10^7 M_\odot$ (e.g. Zwaan et al. 2005; Haynes et al. 2011). However, the redshift evolution of the HIMF remains poorly characterised, awaiting the next generation of facilities. Here we compare our simulated HIMFs with observations, to understand what constraints can be placed on wind models and how the HIMF is expected to evolve.

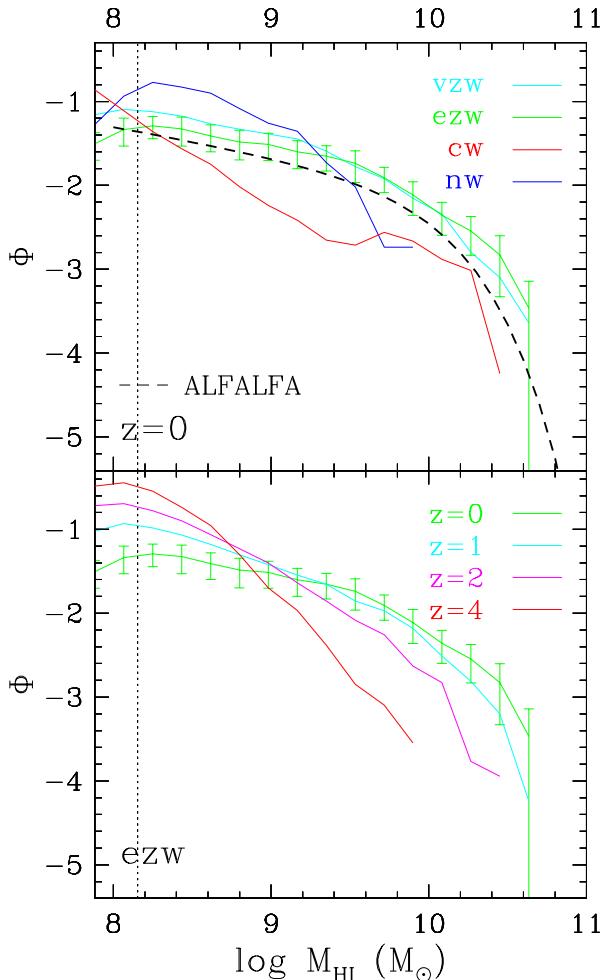


Figure 3. *Top:* A comparison of the HIMF at $z = 0$ in our four wind simulations – ezw (green), vzw (cyan), cw (red), and nw (blue). The vertical dotted line is the stellar mass resolution limit of $1.4 \times 10^8 M_{\odot}$; this is roughly the HI mass resolution limit as well since M_* and M_{HI} are comparable at the smallest masses (see Figure 4). The dashed line shows the HIMF from the $\alpha.40$ sample of the ALFALFA survey (Haynes et al. 2011) for comparison. The ezw and vzw model provide reasonable fits to the data. *Bottom:* Evolution of the HI mass function from $z = 4 \rightarrow 0$ in the ezw model. The low mass end slope increases substantially at high redshifts.

Figure 3, top panel, shows the HIMF for our four wind models. The vertical dotted line shows the *stellar* mass resolution limit of our fiducial r32n512ezw simulation. We have no formal HI mass resolution limit, but since, as we will show in Figure 4, the HI and stellar masses are similar for these small galaxies, this can roughly be considered as the HI mass resolution limit as well. Error bars (shown for ezw) depict cosmic variance as computed from the error of the mean HIMF in each of eight octants of the simulation volume.

For comparison, we show the observational determination of the HIMF from Haynes et al. (2011) (the first 40% of the ALFALFA Survey; dashed line). We note that the ALFALFA observations select by HI mass, while we are effectively selecting by stellar mass; however, the sensitivity of ALFALFA is well below what we can resolve, and hence un-

less there is a dominant population of small HI-free objects (which we will show later does not occur in our models), this comparison should be robust.

Concentrating on the low mass end of the HIMF, we see that both the ezw and vzw models do a good job of matching the observed low mass end slope of $\alpha \approx -1.34 \pm 0.02$ (for the “whole $\alpha.40$ ” sample, from Table 6 of Haynes et al. 2011). In detail, the predicted best-fit Schechter function slope for the ezw model is $\alpha = -1.31$, while for the vzw model it is $\alpha = -1.45$ and, therefore, the ezw model provides a slightly better fit. The massive end is relatively unaffected by quenching, as seen from the minor differences between the ezw model that includes quenching and the vzw that does not include it.

Particularly remarkable is that the amplitudes of the HIMF in these two models are reasonably close to that observed. They are slightly above the observations at all masses, particularly at higher masses. We could in principle tune our values of $N_{\text{HI},\text{lim}}$ and/or shrink the radial extent to which we associate HI with a galaxy within reasonable uncertainties to match the data better, but we prefer to use well-motivated values for these quantities. We have in fact tried several other reasonable values for these quantities, and the net effect is generally to scale the HIMF in amplitude without changing the shape significantly.

The agreement with the low mass end slope of the HIMF, while simultaneously matching the low mass end of the galactic stellar mass function, is a stringent constraint that almost all galaxy formation models have difficulty matching (Mo et al. 2005; Lu et al. 2012, 2013). Fundamentally, in most models this arises because the low mass end slope of the dark matter mass function is quite steep, which is exacerbated by low-mass galaxies being more HI-rich. In our simulations, the stronger outflows from low-mass galaxies strongly suppress the overall baryon content of galaxies. Even more critical is the impact of wind recycling. As described in Oppenheimer et al. (2010), wind recycling is preferentially stronger in high-mass galaxies because the denser surrounding gas slows outflows more effectively via ram pressure. This results in a higher fraction of ejected gas being recycled back into higher-mass galaxies, thereby yielding more star formation in more massive systems. Oppenheimer et al. (2010) showed that this flattens the low mass end slope of the stellar mass function (see Figure 2), and here we see that this effect is also important for the low mass end of the HI mass function. We will discuss further the differences between previous semi-analytic galaxy formation model results and our results in §9.

Both the constant wind (cw) and no wind (nw) models provide significantly poorer fits to the observations than ezw or vzw. Without winds, there is a surplus of low- M_{HI} galaxies, and a deficit of high- M_{HI} systems. This is characteristic of the HIMF in many semi-analytic galaxy formation models (Lu et al. 2012). The former discrepancy arises because there are no outflows to eject material from low-mass galaxies, and hence the baryon fraction in these systems is very large (Davé 2009), in disagreement with observations that show a small baryon content in dwarfs (e.g. McGaugh et al. 2010). The latter arises because, without winds, gas becomes very dense in high-mass galaxies, which means that most of the gas is in molecular form, and much of it has been converted into stars. As discussed in

Davé, Oppenheimer, & Finlator (2011), the stellar content of galaxies in the no-wind model grossly exceeds that observed, and Figure 2 showed that it overproduces the number density of galaxies at all the stellar masses probed.

For the constant wind case, the low mass end slope is quite steep ($\alpha = -1.64$). Moreover, there is a characteristic ‘bump’ in the mass function at $M_{\text{HI}} \sim 10^{10} M_{\odot}$. This bump is reminiscent of a similar bump in the stellar mass function in this model (Davé, Oppenheimer, & Finlator 2011), which arises because wind recycling rapidly becomes important around this mass scale as the (constant velocity) winds are no longer able to escape the galaxy halo’s potential well. Such a feature, which is generic to wind models that assume a constant outflow speed, is not observed in the HIMF (nor in the stellar mass function).

Duffy et al. (2012) employed an outflow model that is similar to this constant wind case. They find that the low mass end slope is fairly flat down to $M_{\text{HI}} \approx 10^9 M_{\odot}$, and was too shallow compared to data. In a comparable range, our HIMF also appears fairly flat – formally, the best-fit low mass end slope ignoring points below that mass is $\alpha = -1.12$, although a Schechter function is a poor descriptor. Our higher-resolution simulation probes further down the mass function, which enables us to see the steep low mass end. Overall, our results for this wind model agree well with Duffy et al. (2012) in the overlapping mass range, and both show that the constant wind model fails to match the low mass end of the HIMF.

The bottom panel shows the redshift evolution of the HIMF out to $z=4$, focusing on our ezw simulation. The low mass end slope becomes progressively steeper with redshift, mimicking the behaviour seen in the stellar mass function (Davé, Oppenheimer, & Finlator 2011). The low mass end slopes at $z = 1, 2, 3, 4, 5$ are $-1.54, -1.79, -1.82, -1.99, -2.11$, respectively ($z = 3, 5$ are not shown). This arises because wind recycling becomes increasingly effective to lower redshifts (Oppenheimer et al. 2010), since the recycling time is roughly constant at $\sim 1 - 3$ Gyr (Oppenheimer & Davé 2008) for $M_{*} \sim 10^{10} M_{\odot}$ galaxies, which is a small fraction of the Hubble time today but comparable to the Hubble time in the early Universe. Hence the HIMF steepens rapidly out to $z \sim 2$, and then the steepening becomes more gradual, since at early times the effect of wind recycling, which is responsible for flattening the HIMF, is reduced. Meanwhile at the massive end, there are fewer galaxies at high redshifts simply because of the hierarchical nature of galaxy assembly.

In summary, the HIMF provides a strong constraint on outflow models. The agreement of the ezw model predictions with the latest observed HIMF from ALFALFA represents a non-trivial success that has not previously been attained in hierarchical models of galaxy formation. The vzw model fares slightly worse but may still be within the overall uncertainties, while the constant wind and no wind cases fare poorly against data. More broadly, this indicates that including a well-motivated model for galactic outflows enables hierarchical structure formation models to produce an HIMF that is in very good agreement with data down to fairly low H I masses.

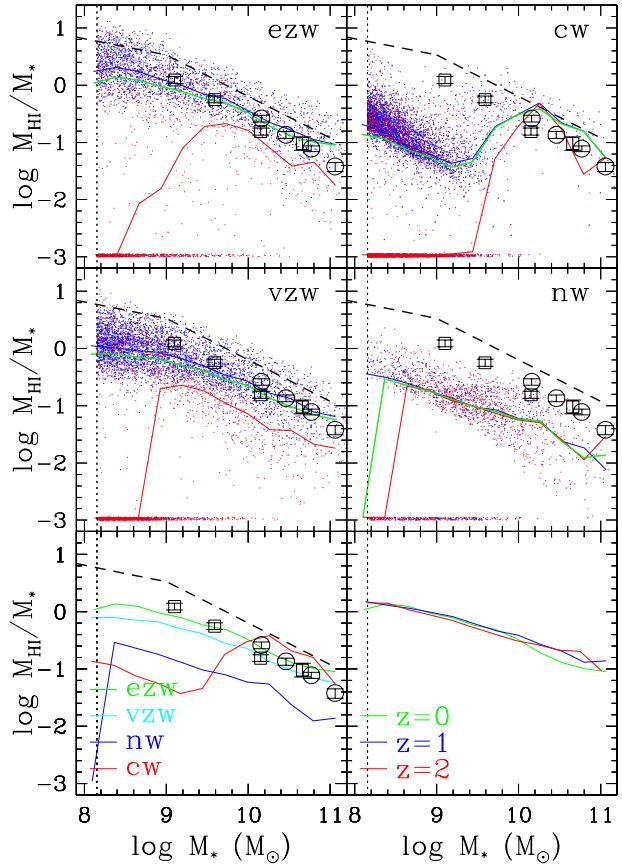


Figure 4. H I richness ($\equiv M_{\text{HI}}/M_{*}$) in galaxies from our four wind models: ezw (upper left), vzw (middle left), cw (upper right), and nw (middle right). Points show individual galaxies at $z = 0$, with central galaxies shown in blue and satellites shown in red. A running median for all galaxies is shown in green, while medians for the centrals and satellites are shown in blue and red, respectively. Circles show H I fractions observed in the GASS survey of nearby galaxies (Catinella et al. 2010), squares show results from the Herschel Redshift Survey (Cortese et al. 2011), and the dashed line shows an approximate fit to the locus traced by the H I -selected ALFALFA survey (Huang et al. 2012). The bottom left panel shows a comparison of all four wind models with observations at $z = 0$, while the bottom right panel shows the evolution of the median H I richness with redshift out to $z = 2$ in the ezw model.

4 HI RICHNESS

The H I richness, i.e. the H I mass relative to the stellar mass, provides a complementary characterisation of the H I content of galaxies. As we argued earlier, both the H I and stellar masses are reasonably robust predictions of our models (unlike the molecular gas content), and hence the H I richness should provide a meaningful discriminant between models.

Figure 4 shows the H I richness in our simulated galaxies, as a function of stellar mass. The top four panels show our four wind models: ezw (upper left), vzw (middle left), constant (upper right), and no winds (middle right). The points in each panel show individual galaxies at $z = 0$, with blue depicting central galaxies and red depicting satellites. Galaxies with zero H I content are plotted along the bottom

of each panel; they are almost exclusively satellite galaxies, which we will examine further in §5. A binned median of $\log M_{\text{HI}}/M_*$ for all galaxies is shown as the green line. We show errorbars on the median corresponding to the 1σ spread for the galaxies within each mass bin. We also separately show binned medians for the central and satellite populations. The vertical dotted line shows our galaxy stellar mass resolution limit.

For comparison, mean M_{HI}/M_* observations from the GALEX Arecibo SDSS Survey (GAASS) are shown as the large circles (Catinella et al. 2012). GAASS has measured HI in a stellar mass-limited sample down to $M_* \sim 10^{10} M_\odot$. At lower masses, we show the results from the Herschel Reference Survey (HRS) by Cortese et al. (2011, open squares), which uses literature HI data but is approximately stellar mass-complete. We also show the fit to results from the ALFALFA survey by Huang et al. (2012, dashed line); since this is HI-mass selected, it is biased towards more HI-rich galaxies, as is evident when compared to the M_* -selected data. Our simulation results are most straightforwardly comparable to M_* -limited samples.

All wind models broadly predict that HI richness is anti-correlated with M_* , as observed. However, in detail, the models show distinct differences; for clarity, just the overall medians are plotted in the lower left panel. Our ezw model produces a good agreement with the stellar mass-selected observations down to the lowest probed masses. The trend for the vzw model follows ezw, but shows slightly lowered HI richness particularly in smaller galaxies, indicating that more mass-loaded galactic outflows in low-mass galaxies are favoured. The no-wind case follows the trend of the data, but is too low by a factor of a few in HI richness, showing that there is overly efficient conversion of gas into stars in this model. Finally, the constant wind model produces a distinct feature in HI richness, mimicking the feature seen in the HIMF, owing to wind recycling. This agrees poorly with the observations, which has no comparable feature. Note that more highly mass-loaded outflows yield *higher* HI richness, not lower as one might naively expect, in part because such outflows suppress M_* .

The redshift evolution of HI richness is shown in the lower right panel of Figure 4, for our ezw simulation. Despite the rapid evolution in the HIMF, there is remarkably little evolution in the HI fraction at a given stellar mass from $z = 0 \rightarrow 2$. The slow evolution in HI content is a contrast to the much more rapid evolution observed in the molecular gas fraction (e.g. Tacconi et al. 2010; Geach et al. 2011), although some of that evolution may reflect uncertainties in assessing the molecular gas content from CO emission (e.g. Narayanan, Bothwell, & Davé 2012; Bolatto, Wolfire, & Leroy 2013). Also, the lack of evolution in HI richness indicates that the evolution in the HIMF discussed in §3 mostly reflects the evolution in the stellar and/or halo mass functions. We will see in §7 that the lack of evolution in global HI content is also a generic prediction of our models.

5 HI IN SATELLITE GALAXIES

The HI content of galaxies is seen to vary substantially with environment. In the densest environs such as clus-

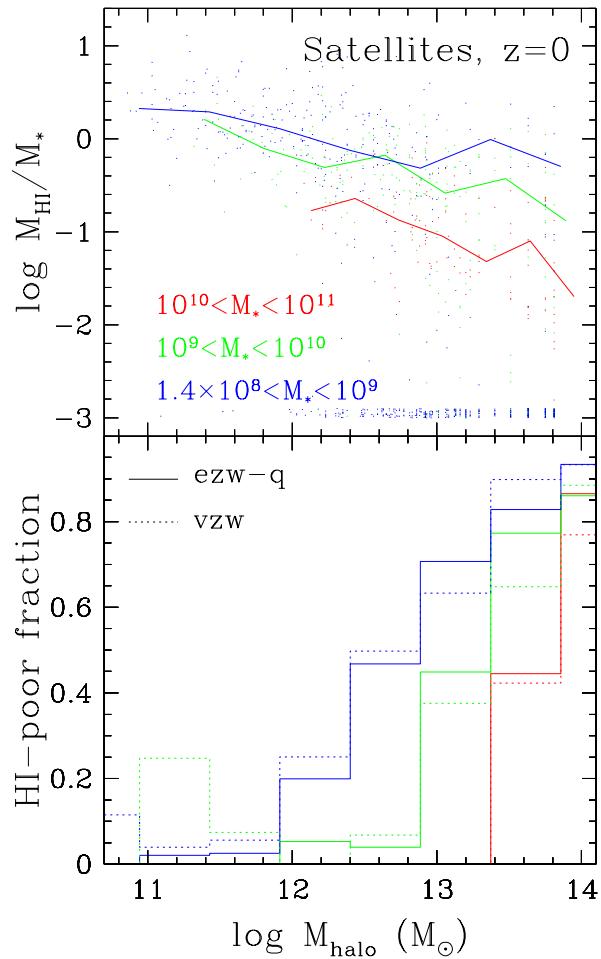


Figure 5. The top panel shows HI richness ($\equiv M_{\text{HI}}/M_*$) for satellite galaxies versus halo mass at $z = 0$, for our ezw model. The satellites are split into low (blue), intermediate (green), and high (red) stellar mass bins, as indicated. Satellites with no HI are plotted at -3 , with an artificial spread to aid visualisation. Lines show a running median of only the galaxies that have HI (i.e. ignoring the points plotted along the bottom). The bottom panel plots the fraction of HI-free satellites within those same mass bins, colour coded as above. Solid lines show results from the ezw model that includes an ad hoc quenching model at high halo masses, while dotted lines show the equivalent plots for the vzw model that does not include such quenching.

ters, ram pressure can remove HI from infalling galaxies, and it is long known that cluster galaxies are deficient in HI (e.g. Haynes et al. 1984). Even in less extreme environments where ram pressure stripping is expected to play less of a role, the HI content appears to be anti-correlated with local density (Robertson, Shields, & Blanc 2012); other processes such as strangulation or harassment may also be playing a significant role. In general, such environmental processes are expected to preferentially lower the HI content of satellite galaxies, particularly those within larger halos that are expected to have hotter ambient gas. In this section we examine how the HI content of satellites varies with halo mass, which can be considered as a rough proxy for environment.

Returning to Figure 4, it is clear that satellites (red

points and curve) have a lower $\text{H}\alpha$ richness than central galaxies (blue) of the same stellar mass. At the high mass end, the typical difference in median $\text{H}\alpha$ richness between central and satellites is a factor of $\sim 2 - 3$, in all the wind models. Below $M_* \lesssim 10^{9.5} M_\odot$, the satellite $\text{H}\alpha$ fraction drops very quickly, as there are numerous low-mass $\text{H}\alpha$ -free satellites in our simulations (plotted along the bottoms of the panels). While the majority of galaxies of low mass are still centrals (Davé, Oppenheimer, & Finlator 2011), and hence the overall median tracks that of central galaxies, it is clear that low-mass satellites in particular are highly deficient in $\text{H}\alpha$ for their mass.

Figure 5 examines these trends more closely. The top panel shows the $\text{H}\alpha$ richness in satellite galaxies in our ezw simulation at $z = 0$, divided into three stellar mass bins of $1.4 \times 10^8 - 10^9 M_\odot$, $10^9 - 10^{10} M_\odot$, and $> 10^{10} M_\odot$. Galaxies with $\text{H}\alpha$ richness less than 10^{-3} (virtually all of which are $\text{H}\alpha$ -free) are plotted along the bottom at -3 ; lines show a running median *not* including these $\text{H}\alpha$ -free galaxies.

This plot shows several key trends. First, at all masses, the satellites that have $\text{H}\alpha$ show a mild trend of being more $\text{H}\alpha$ -rich in lower mass halos. Since the halo mass traces stellar mass in our models (Davé, Oppenheimer, & Finlator 2011), this basically reflects the fact that lower stellar mass galaxies have higher $\text{H}\alpha$ richness (as seen in Figure 4). Similarly, lower-mass satellites are more $\text{H}\alpha$ rich, again following the trend for centrals discussed in the previous section. Overall these trends appear to reflect the $\text{H}\alpha$ content of the satellite galaxies when they were still centrals.

However, a clear difference between central and satellite galaxies is that there are many more satellites that are devoid of $\text{H}\alpha$, particularly at higher halo masses, i.e. the points along the bottom of the plot appear much more frequently in the more massive halos. In fact, the distribution of $\text{H}\alpha$ richness in satellites appears to be bimodal. Even in large halos, there are still some satellites that have substantial $\text{H}\alpha$, and their $\text{H}\alpha$ content is not grossly different than that of satellites in smaller halos. However, the *fraction* of satellites that are $\text{H}\alpha$ poor increases sharply in higher mass halos.

To quantify this, we plot in the bottom panel of Figure 5 the fraction of galaxies with $\text{H}\alpha$ richness $< 10^{-3}$, as a function of halo mass, divided again into the three satellite galaxy stellar mass bins as in the top panel. Here, we clearly see that at any halo mass, lower-mass satellites are more likely to have had their $\text{H}\alpha$ content strongly reduced. Moreover, this trend is a very strong function of halo mass, with satellites of all masses being much more likely to be $\text{H}\alpha$ -poor if they lie within a more massive halo.

The strong bimodality suggests that the process that renders satellites $\text{H}\alpha$ poor happens on a relatively short timescale compared to the infall timescale into the halo (which is roughly comparable to the halo's dynamical time of several Gyr). In future work we plan to investigate the detailed dynamical processes that remove the $\text{H}\alpha$ from satellites in our simulations.

The dotted lines in the bottom panel of Figure 5 show the analogous results for our vzw simulation. In massive halos where the fraction of $\text{H}\alpha$ -poor satellites is substantial, the outflow model is identical between vzw and ezw simulations, but the ezw simulation includes our quenching prescription. The fact that the dotted and solid lines are very similar indicate that the trends seen in the $\text{H}\alpha$ content of satel-

lites are not being set by our ad hoc quenching prescription, but more likely by the fact that halo masses above $10^{12} M_\odot$ tend to contain much more hot halo gas that can more strongly impact satellites moving through it (Kereš et al. 2005; Gabor et al. 2011). Recall that our quenching prescription applies only to galaxies with a high velocity dispersion, which is typically only the central galaxy in the halo, and hence the only direct impact on smaller satellites would be from the extra heat being added to the halo gas; evidently this has a minimal effect on the satellites in our simulations.

What about the central galaxies? In general, very few of the central galaxies are devoid of $\text{H}\alpha$, as seen in Figure 4. However, although we don't highlight it, there exists a small population of lower-mass ($M_* \lesssim 10^{10} M_\odot$) centrals that are $\text{H}\alpha$ -poor. This is related to what was seen by Gabor et al. (2012), who found numerous low-mass central galaxies on the red sequence. These turned out to be galaxies that reside just outside, i.e. within several virial radii, of more massive halos. In the spherical overdensity algorithm we use to identify halos, such galaxies are identified as centrals, although the influence of the larger galaxy's halo can extend to well beyond its virial radius (e.g. Moldar et al. 2009). Hence these $\text{H}\alpha$ -poor centrals could be impacted by the extended environment of a nearby larger halo, or else they could be former satellites whose orbit has taken them outside the nominal virial radius.

In summary, halo mass plays an increasingly important role in setting the $\text{H}\alpha$ content of satellite galaxies. This is particularly seen by the strongly increasing fraction of $\text{H}\alpha$ -poor satellites as a function of halo mass. At a given halo mass, low-mass satellites have a greater chance of having their $\text{H}\alpha$ removed. Satellites that have not had most of their $\text{H}\alpha$ removed lie along similar relations to satellites in lower-mass halos. These results suggest that the process by which $\text{H}\alpha$ is removed from satellites in our simulations acts fairly quickly, and preferentially on smaller galaxies. Comparing these predictions to observations can help constrain such $\text{H}\alpha$ removal mechanisms.

6 $\text{H}\alpha$ DEFICIENCY

The $\text{H}\alpha$ richness of galaxies is observed to correlate with a variety of galaxy properties (besides stellar mass). In the previous section we showed that satellite galaxies in higher mass halos were increasingly stripped of their $\text{H}\alpha$. But even galaxies that still have substantial $\text{H}\alpha$ show correlations of their $\text{H}\alpha$ content with properties such as environment (e.g. Cortese et al. 2011), and metallicity (Robertson, Shields, & Blanc 2012). Here we examine the second-parameter trends of $\text{H}\alpha$ with environment, star formation rate, and metallicity in our simulations to provide insights into the physical drivers that establish the $\text{H}\alpha$ content of galaxies.

To set a theoretical context for this discussion, we recall the equilibrium model for galaxy growth as described in Davé, Finlator, & Oppenheimer (2012), which presents a simple physical scenario that can account for the relationships between key galaxy properties, as well as the scatter around those relationships. In the equilibrium model, accretion onto a galaxy from the cosmic web is fairly quickly

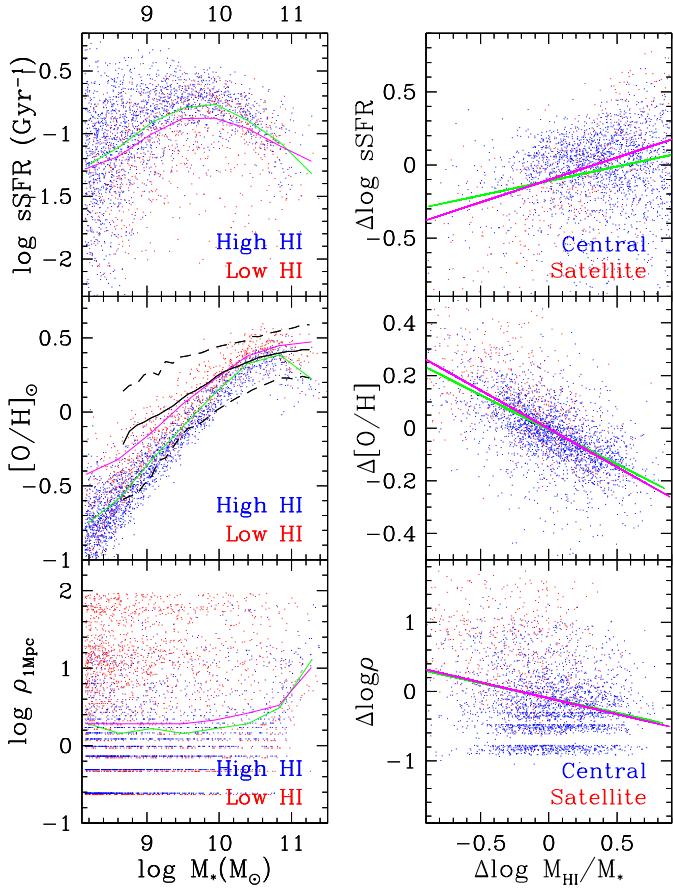


Figure 6. *Left panels:* Specific SFR, metallicity, and environment vs. stellar mass (top to bottom) from our ezw simulation at $z = 0$, with the blue and red points indicating galaxies that have above and below (respectively) the median $\text{H}\alpha$ richness at that given M_* . The green line shows the running median. The cyan line shows the median for the vzw model, although the points are not shown. The black solid line in the middle-left panel plots the median observed relation from (Tremonti et al. 2004) with the dashed lines indicating the 95% contours. $\text{H}\alpha$ -rich galaxies at a given M_* tend to have higher sSFR, lower metallicity, and live in less dense regions. *Right panels:* The deviation from the median sSFR, metallicity, and environment (top to bottom) at a given M_* vs. the deviation from the median $\text{H}\alpha$ richness. The magenta lines show power law fits to the deviations, having slopes of 0.31, -0.26, and -0.56, respectively.

processed into either stars or an outflow, resulting in a slowly-evolving gas reservoir. This results in fairly tight relations between stellar mass and star formation rate (the so-called main sequence, e.g. Davé 2008), gas-phase metallicity (Finlator & Davé 2008), and star-forming gas content (Davé, Finlator, & Oppenheimer 2011). The predicted relations are in good agreement with the hydrodynamic simulation results.

Stochasticity in the inflow rate causes scatter about these equilibrium relations. As described in Davé, Finlator, & Oppenheimer (2011), an accretion event such as a minor merger will cause an increase in gas content that raises the star formation rate, while simultaneously lowering the metallicity. Conversely, a lull in ac-

cretion, or diminished accretion owing to a galaxy becoming a satellite in a larger halo, will cause the existing gas to be consumed, resulting in a lower gas content, a lower SFR, and a higher metallicity. Hence, for example, departures from the mass-metallicity relation are expected to be inversely correlated with SFR. Such a trend has been observed (Lara-López et al. 2010; Mannucci et al. 2010), and is known as the fundamental metallicity relation; our simulations naturally predict this (Davé, Finlator, & Oppenheimer 2011). Our models analogously predict that deviations in the star-forming (i.e. mostly molecular) gas content will inversely correlate with deviations in metallicity (Davé, Finlator, & Oppenheimer 2011).

We now extend this to consider the impact of inflow stochasticity on the $\text{H}\alpha$ reservoir of galaxies. The left panels of Figure 6 show the specific SFR, the metallicity, and the local galaxy density averaged over 1 Mpc spheres as a function of stellar mass in our ezw simulation at $z = 0$. Metallicity here is computed as the SFR-weighted oxygen metallicity (Finlator & Davé 2008; Davé, Finlator, & Oppenheimer 2011). The observed mass-metallicity relation (MZR) from SDSS (Tremonti et al. 2004) is shown as the solid black line with enclosing 95% contours indicated by the dashed lines. We have converted these metallicities to solar units assuming a solar oxygen abundance from Asplund et al. (2009), namely $12 + \log [\text{O}/\text{H}]_{\odot} = 8.70$. The green line shows a running median in each panel. We also show as the magenta line the median for the vzw simulation (without showing the individual points), for comparison.

The median sSFR drops slowly with M_* at both small and large stellar masses, with a peak at $M_* \sim 5 \times 10^9 M_{\odot}$; there are minimal differences between the ezw and vzw models. The vzw line follows that in Davé, Oppenheimer, & Finlator (2011), who used lower-resolution simulations that only probed down to $M_* \sim 10^9 M_{\odot}$, while here the turnover is much more apparent. We note that SDSS observations do not indicate such a turnover (Salim et al. 2007), which reflects a generic problem in hierarchical models that star formation in dwarfs peaks too early and is too low today (Weinmann et al. 2012).

The metallicity is tightly correlated and increases with M_* until $M_* \sim 5 \times 10^{10} M_{\odot}$. The ezw model produces a steeper mass-metallicity relation (MZR) than the vzw model, for reasons discussed in Finlator & Davé (2008) and Davé, Finlator, & Oppenheimer (2012). The metallicity $Z \propto 1/(1 + \eta)$ (in the absence of wind recycling), with $\eta \propto 1/\sigma \propto M_*^{-1/3}$ for all galaxies in the vzw model and $\eta \propto 1/\sigma^2 \propto M_*^{-2/3}$ for the ezw model at small masses. For $\eta \gg 1$ as at small masses, this toy model yields $Z \propto M_*^{1/3}$ for the vzw model and $Z \propto M_*^{2/3}$ for the ezw model. This represents an asymptotic slope, and since the ezw model follows the vzw model scalings at high masses, the actual MZR in the ezw model is not as steep as this, but it is clearly steeper than in the vzw model. The vzw model appears to provide a better fit to the SDSS data, although metallicity measures are sufficiently uncertain that the ezw prediction cannot be ruled out (Ellison et al. 2008). Also, a recent determination of the MZR from direct metallicity measures by Andrews & Martini (2013) suggests a steeper MZR slope at small masses, closer to $Z \propto M_*^{1/2}$.

The simulated galaxies in Figure 6 are colour-coded by H_I richness: the blue points show galaxies that have a higher-than-median H_I richness at that stellar mass, and the red points show the converse. It is clear that H_I richness has a strong second parameter correlation with all these properties: H_I-rich galaxies at a given mass also tend to have a higher sSFR, a lower metallicity, and a lower local galaxy density.

To reiterate the physical scenario: When accretion happens, it results in galaxies with higher H_I richness than typical. This high H_I richness also corresponds to a lower-than-normal gas-phase metallicity. This suggests that the H_I content of galaxies can be an indicator of recent accretion, as argued from observations by Moran et al. (2012), who show that the low metallicities in H_I-rich systems most strongly appear at the outskirts of disks.

We can quantify these trends using deviation plots, which we show in the corresponding right panels of Figure 6. This shows the difference between a galaxy's sSFR, metallicity, and density with the median value of these quantities at a given M_* , i.e. approximately the green line in the left panels (although in detail we employ a spline fit to produce more smoothly varying deviations), versus the deviation in H_I richness with its median value at that M_* . Points in these panels are colour-coded by centrals (blue) and satellites (red). The green lines show power-law fits to these deviations for the ezw model while the magenta lines show the same for the vzw model (though the individual points are not shown).

These right panels quantify the second-parameter trends from the left panels: galaxies with high H_I content also have a high sSFR, a low metallicity, and are in low density environments. A power law provides a reasonable fit to the sSFR and metallicity data, with slopes of 0.31 and -0.26 , respectively. The vzw model yields virtually identical slopes, even though the trends (such as the MZR) are noticeably different. This occurs because the trend in the scatter arises from inflow stochasticity, which has little to do with the outflows. The central and satellite galaxies do not show significantly different trends in sSFR or metallicity.

The metallicity dependence can be compared to recent work by Robertson, Shields, & Blanc (2012), who looked at the H_I deficiency parameter, *DEF*, relative to the deviation from the expected oxygen abundance for cluster and field galaxies, finding slopes of -0.25 ± 0.12 and -0.41 ± 0.14 , respectively. Our simulated galaxies are more comparable to field galaxies, which show a higher slope but are still within their uncertainties. Robertson, Shields, & Blanc (2012) compared to results from the simulations in Davé, Finlator, & Oppenheimer (2011), which yielded a similar slope. Although the definition of DEF differs from the deviation we plot here, the qualitative agreement in the trend indicates that the basic physical model of a slowly-evolving equilibrium in gas content and metallicity appears to be broadly consistent with the observations.

Finally, we consider the environment in the bottom two panels of Figure 6. The median local galaxy density is independent of mass until the very highest mass systems in our volume at $M_* \gtrsim 5 \times 10^{10} M_\odot$. Unlike in the sSFR and metallicity deviation plots, where the centrals and satellites do not significantly deviate from one other, the envi-

ronment deviation plot (bottom right) markedly separates the central and satellite galaxies. This environmental dependence arises because the environment can impact the inflow rate into satellites. High density regions associated with halos with masses $\gtrsim 10^{12} M_\odot$ will contain substantial amounts of hot gas (Kereš et al. 2005, 2009; Gabor et al. 2012) that can retard accretion (Dekel & Birnboim 2006; Davé, Finlator, & Oppenheimer 2012). Hence galaxies in such regions, particularly satellites, will have lower accretion rates compared to field galaxies at the same mass, and hence less H_I. Overall, the full galaxy sample trends towards having less gas-rich galaxies in denser environments, but this is driven almost entirely by the satellites. Fitting a power-law formally yields a slope of -0.56 (shown as the magenta line), though a power law does not appear to be a particularly good fit.

In summary, the correlated deviations between sSFR, metallicity, and gas content provide quantitative constraints in the way in which galaxies oscillate about their equilibrium relations owing to stochastic accretion. The resulting deviation slope reflects the correlation between the instantaneous gas inflow rate, its conversion to neutral hydrogen, and the infalling metallicity, as mediated by the local environment. Such constraints provide valuable discriminants between different physical scenarios for galaxy growth. The consistency with available observations, albeit preliminary and qualitative, indicates that the stochasticity in the inflow as expected from cosmological accretion is able to broadly explain the second parameter trends observed in the relationship between H_I content and other physical galaxy properties.

7 Ω_{HI} EVOLUTION

Most neutral hydrogen in the cosmos exists in and around galaxies, between the regime where self-shielding happens in the outskirts of disks, to where gas becomes molecular-dominated within star-forming regions. While it is currently infeasible to probe the evolution of H_I in galaxies via 21cm emission to high redshifts³, other avenues have been employed to measure cosmic H_I evolution, such as the abundance of damped Ly α systems (DLAs). In this section we examine the evolution of the cosmic H_I density in our simulations directly from the galaxy population, and compare it with various observational determinations.

Figure 7 shows the evolution of Ω_{HI} from $z = 5 \rightarrow 0$ in our vzw model, for all galaxies above our approximate H_I mass resolution limit of $M_{\text{HI}} > 1.4 \times 10^8 M_\odot$ (solid blue line), as well as for galaxies with $M_{\text{HI}} > 10^9 M_\odot$ (dashed) and $M_{\text{HI}} > 10^{10} M_\odot$ (dotted). For comparison, we also show the evolution of the stellar mass density Ω_* as the red line. Observations of Ω_{HI} at $z > 2$ are shown from Prochaska & Wolfe (2009) derived from DLA abundances (circles with errorbars). At $0.5 < z < 2$, we show a box that encompasses the 1σ range from Rao et al. (2006), who used strong Mg II systems as a proxy for DLAs. Finally, the triangle at $z = 0$ shows Ω_{HI} measured from ALFALFA (Haynes et al. 2011). Note that our prescription for

³ This is a key goal of the upcoming LADUMA survey using the MeerKAT array (Holwerda et al. 2011)

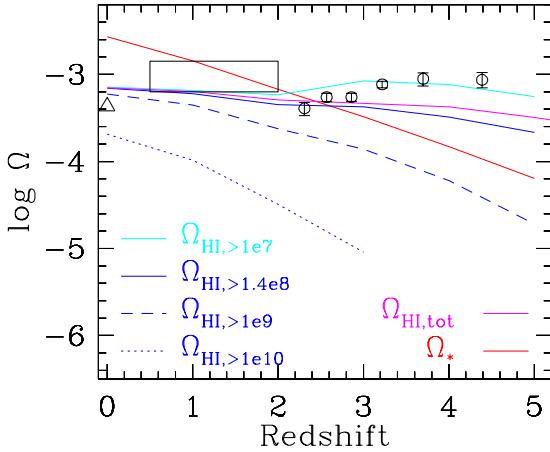


Figure 7. Ω_{HI} , the fraction of cosmic mass density in H I, from $z = 5 \rightarrow 0$ from our ezw simulation. The blue solid line shows Ω_{HI} from all resolved galaxies ($M_* > 1.4 \times 10^8 M_\odot$) the blue dashed and dotted lines show Ω_{HI} in galaxies with $M_* > 10^9 M_\odot$ and $M_* > 10^{10} M_\odot$, respectively. The cyan line shows Ω_{HI} from extrapolating a Schechter function fit to the HIMF at each redshift down to $M_{\text{HI}} = 10^7 M_\odot$. Magenta line shows the total H I mass density in the entire volume. For comparison, the red line shows the mass density in stars in this model. Observations are shown at $2 \lesssim z \lesssim 4.5$ (circles with errorbars) from Prochaska & Wolfe (2009), at $0.5 \lesssim z \lesssim 2$ (box encompassing 1σ limits) from Rao et al. (2006), and at $z = 0$ (triangle) from Haynes et al. (2011).

determining self-shielding calibrated to radiative line transfer simulations as described in §2 does reasonably well predicting Ω_{HI} at $z = 0$; we do not adjust any parameters to match this data as was done in Popping et al. (2009).

Our simulations predict that Ω_{HI} is remarkably constant from $z = 5 \rightarrow 0$, essentially unchanging to within 50% over these 12 billion years. This is in contrast to the dramatic increase of Ω_* over that time interval; over 98% of all stars in this model form since $z = 4$, with more than 80% since $z = 2$ (in broad agreement with recent observational estimates; e.g. Leitner 2012). It is also quite different than the rapid evolution in cosmic star formation rate density (e.g. Hopkins & Beacom 2006; Fardal et al. 2007) and molecular gas fraction (Tacconi et al. 2010; Geach et al. 2011; Tacconi et al. 2012), which both increase by an order of magnitude or so out to $z = 2$ (though see Narayanan, Bothwell, & Davé 2012, who argue for less growth in the observations owing to variations in the CO-to-H₂ conversion factor). This emphasises that H I represents a transient reservoir in the journey of gas from the ionised IGM to stars forming deep within galaxies, and is not directly proportional to the amount of stars being formed or molecular gas present. This cautions against over-interpreting quantities such as the “H I star formation efficiency”, i.e. the SFR divided by the H I mass, since the two quantities are not directly related.

Subdividing Ω_{HI} into different H I mass bins, we see that at $z = 0$, most of the cosmic H I is in galaxies with $10^9 < M_{\text{HI}} < 10^{10} M_\odot$. Going back in time, H I shifts towards lower mass galaxies; by $z = 2$, only half the H I is in $M_{\text{HI}} > 10^9 M_\odot$ systems. In contrast, only a very

small portion of the cosmic H I ever resides in galaxies with $M_{\text{HI}} > 10^{10} M_\odot$, since higher mass galaxies quickly drop off in their H I richness (Figure 4). Thus quenching mechanisms, which primarily affect high-mass galaxies (e.g. Gabor et al. 2012), are expected to have almost no effect on Ω_{HI} . We also show, as the magenta line, the total H I mass density in the entire simulation volume. The difference between the solid magenta and blue lines reflects H I that is not within resolved galaxies (including the diffuse IGM). This extra contribution is a couple percent at low redshift but rises to $\sim 50\%$ at $z = 5$, which again reflects the increasing contribution of small (unresolved) galaxies to the global H I budget at high redshifts.

Observations of Ω_{HI} likewise indicate very little evolution from $z \sim 4 \rightarrow 0$. Different observational tracers generally agree to within a factor of approximately 2, and together indicate essentially no change in Ω_{HI} for the past 12 billion years, in broad agreement with our predictions. Our simulation begins to underpredict Ω_{HI} at $z \gtrsim 3$ when compared with the DLA data. The discrepancy could be physical or numerical. A physical explanation could be that lower metallicity galaxies at early epochs will tend to have less efficient conversion of their atomic gas to molecular gas owing to lower cooling rates and harder interstellar radiation, and hence our locally-calibrated prescription for R_{mol} from Leroy et al. (2008) may not be appropriate.

A numerical explanation for the discrepancy may be that we do not fully resolve galaxies with $M_* < 1.4 \times 10^8 M_\odot$, and there may be substantial contributions to Ω_{HI} from the very smallest galaxies. At low- z , the low mass end of the HIMF is fairly shallow, so the expected contribution from lower mass galaxies is small. But at higher redshifts, Figure 3 shows that the slope becomes substantially steeper, meaning that the additional contribution from unresolved systems could be large.

We can crudely correct for this by fitting a Schechter function to the HIMF at each redshift as we did in §3, and then integrating down to some chosen lower mass limit. As an illustration, we integrate down to $M_{\text{HI}} = 10^7 M_\odot$, which is around the lowest mass observable at low- z in large surveys such as ALFALFA. The result of this Schechter fit extrapolation to $M_{\text{HI}} = 10^7 M_\odot$ is shown as the cyan line in Figure 7. This results in a negligible correction at low- z , but at higher redshifts the correction can be up to a factor of three, which agrees better with the high- z DLA results. We caution that this exercise is intended to be illustrative, since our M_{HI} limit was chosen rather arbitrarily, and it may not correspond to the effective H I masses probed by DLA systems. Furthermore, our Schechter function fits become increasingly uncertain at higher redshifts, owing to the lack of dynamic range in our simulations. Nonetheless, this illustrates that plausible corrections down to lower H I masses can bring our predicted Ω_{HI} into better agreement with the observations by preferentially increasing the high-redshift H I mass density. Given the crudeness of our prescriptions for determining H I content, agreement at this level is encouraging.

In summary, our simulations generally predict a very slowly-evolving cosmic H I mass density, in broad agreement with the observations. However, we note that the HIMF actually evolves rather considerably (Figure 3), with many more low- M_{HI} galaxies and fewer high- M_{HI} ones at high z .

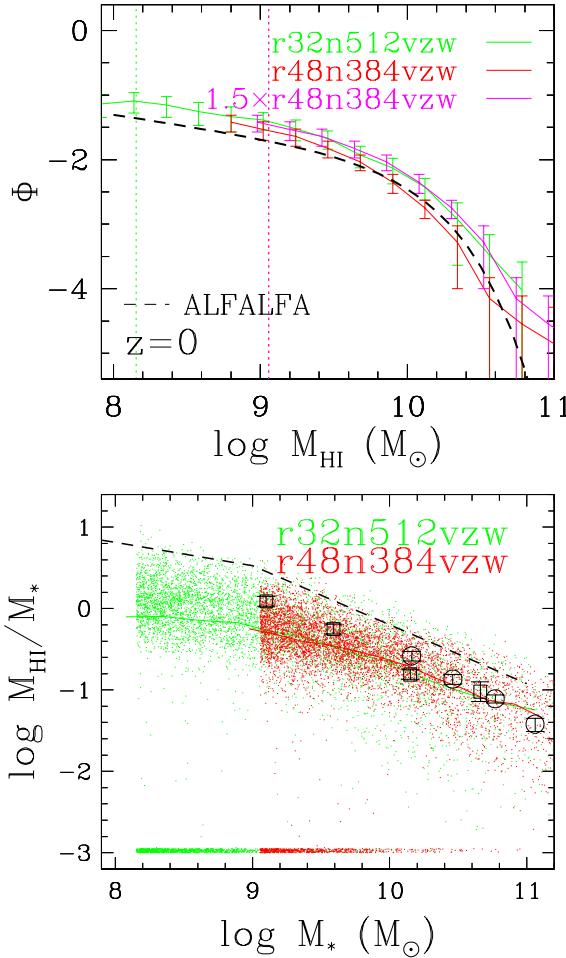


Figure 8. Comparison of the H I mass function (top panel) and H I richness (bottom panel) in our fiducial $32h^{-1}\text{Mpc}$, 2×512^3 particle simulation (green) compared to a $48h^{-1}\text{Mpc}$, 2×384^3 particle simulation (red) with $8\times$ poorer mass resolution. Resolution convergence is good for H I richness, while the HIMF is increased in the higher-resolution simulation, but is still within the formal uncertainties. We show, as the magenta line in the top panel, the r48n384 HIMF where the H I masses (not Φ) have been multiplied by 1.5; this results in a nearly identical mass function to the r32n512 case.

It is something of a coincidence that these two variations roughly cancel when summed to give the global H I mass density. Nonetheless, this emphasises the transient nature of the H I reservoir around galaxies, which does not build up hierarchically in the way that the stellar mass does, but instead responds to evolution in the ionising background and the ISM pressure.

8 RESOLUTION CONVERGENCE

In this study we have employed some of the highest-resolution hydrodynamic simulations of random cosmological volumes ever evolved down to $z = 0$. Nonetheless, it is important to assess whether our resolution is sufficient to robustly predict the H I properties of galaxies. Here we examine our two most basic H I statistics, namely the

H I mass function and the H I richness vs. stellar mass, in a simulation that has 2×384^3 particles in a $48h^{-1}\text{Mpc}$ volume (r48n384), described in our previous work (e.g. Oppenheimer et al. 2010; Davé, Oppenheimer, & Finlator 2011; Davé, Finlator, & Oppenheimer 2011). This simulation has 8 times poorer mass resolution and 2 times poorer spatial resolution, albeit in a volume that is 3.4 times larger. We use the identical momentum-driven wind (vzw) and cooling model in both simulations⁴.

Figure 8 shows the H I mass function (bottom) and the M_{HI}/M_* ratio vs. M_* (top) and in these two simulations, r32n512vzw in green and r48n384vzw in red. These are analogous to Figures 4 and 3, respectively, and the observations are plotted as in those figures.

For the H I fraction, the resolution convergence is excellent; the running median lines for the high and low resolution simulations lie on top of each other in the mass ranges that overlap. It shows that the anti-correlation between H I richness and stellar mass is converged even at significantly poorer mass resolution. Hence, at least over this modest dynamic range in mass resolution, the H I richness of galaxies is a well-converged property in these simulations.

In contrast, the HIMF (top panel) is less well converged. In general, the number density at a given H I mass increases with resolution. As shown by the magenta line, the two mass functions lie on top of each other if the H I mass in each low-resolution galaxy is multiplied by a factor of 1.5. Yet the good agreement with the H I richness plot would suggest that the differences in the HIMF between these models actually reflects differences in the stellar mass function. Indeed this is seen to be true (not shown), which may arise owing to increased wind recycling in higher-resolution simulations, or simply owing to the cosmic variance within the smaller volume; we leave this investigation for the future. Unfortunately, we do not have even higher resolution simulations within a sufficiently large volume to assess whether this converges with increased resolution. In any case, the differences in Φ between the models at any given M_* is smaller than the variance within that bin.

Overall, the resolution convergence for H I properties is good but not ideal. For a given M_* the H I content is very robust, but the number density of galaxies at a given H I mass is less well converged. For that reason, Ω_{HI} is also slightly uncertain, as there is a $\sim 30\%$ difference between the $z = 0$ value predicted by the high and low resolution simulations. We do not yet know whether these differences reflect a resolution effect or a volume effect; we would need a wider suite of simulations to assess these issues. While these discrepancies are not trivial, they do not significantly affect the main conclusions of this paper.

9 SUMMARY AND DISCUSSION

The H I content of galaxies offers a unique glimpse into baryon cycling processes of inflow and outflow that are believed to govern galaxy evolution. Hydrodynamic simulations can provide a way to connect H I to such baryon cycling

⁴ We use the vzw wind model because we already have the r48n384vzw simulation in hand, and there are usually only minor differences between vzw and our now-favoured ezw model.

processes, as well as to help interpret observations within a hierarchical galaxy formation context. In this work, we have examined the HI properties of galaxies in cosmological hydrodynamic simulations with four different outflow prescriptions. Our key conclusions can be summarised as follows:

- The HI mass function is strongly affected by the inclusion of galactic outflows. Without outflows, our simulations yield too many galaxies at small HI masses, and too few at large HI masses. Our model with a constant wind speed (cw) introduces a feature in the HIMF around $M_{\text{HI}} \sim 10^{10} M_{\odot}$ that is not seen in the observations. Our models with momentum-driven wind scalings produce an HIMF that broadly agrees with the observations, particularly at the low mass end down to our mass resolution limit of $\sim 10^8 M_{\odot}$. Switching to energy-driven scalings at low masses ($\sigma_{\text{gal}} < 75$ km/s) further improves the agreement, and generates a low mass end HIMF slope of -1.3 . The HIMF is slightly above the data at all masses, but likely within the systematic uncertainties regarding how we compute HI masses for our galaxies. The low mass end slope becomes progressively steeper with redshift out to $z \sim 3$, then remains constant at around -2.1 at higher redshifts. As an aside, we showed that the ezw model also provides an excellent match, within the quoted uncertainties, to the galaxy stellar mass function.

- All models show that low-mass galaxies are more HI rich, relative to their stellar content, than high mass galaxies. The exact shape and amplitude of the relation depends slightly on the wind model. All the models broadly agree with the GASS data that probes down to $M_* \sim 10^{10} M_{\odot}$, but the different wind models predict different trends to lower masses. The constant wind model predicts significantly lower gas richness than our vzw (momentum-driven wind scalings) or ezw models at $M_* \lesssim 10^{10} M_{\odot}$, reflecting a strong suppression of inflow owing to energetic winds (van de Voort et al. 2010) that likewise suppresses specific SFRs (Davé, Oppenheimer, & Finlator 2011). Current observations favour a continuing increase of gas richness to lower masses, and our ezw model again fares slightly better than vzw at low masses, and both match much better than the cw model. There is essentially no redshift evolution in the HI richness as a function of stellar mass, showing that the evolution of the HIMF is interconnected with the evolution of the galaxy stellar mass function.

- Galaxies with a high HI content tend to have lower gas-phase metallicities and higher star formation rates at a given M_* . The deviation in sSFR and metallicity versus the deviation in HI richness can be fit by power laws with slopes of 0.31 and -0.26 , respectively. The latter slope compares favourably with existing observations of the HI deficiency parameter versus metallicity deviation. These deviation trends can be understood within the context of an equilibrium model in which galaxies tend to have equilibrium relations in star formation rate, metallicity, and gas content, perturbed by stochastic accretion events that result in correlated deviations in SFR, metallicity, star-forming gas mass, and HI mass. This suggests that HI content is an indicator of recent accretion events in the outskirts of galaxies that stimulates star formation.

- Environment plays an important role in governing HI content, causing an increasing suppression of the HI content

of satellites in more massive halos. The suppression is rapid, as indicated by the bimodal distribution of the HI fraction at a given halo mass. Lower mass satellites are more likely to have their HI removed. The deviation of HI with environment is anti-correlated with the deviation in HI richness, driven by HI-poor satellite galaxies in denser regions.

- The global HI mass density evolves slowly from $z \sim 5 \rightarrow 0$, in broad agreement with the observations of DLAs and other measures. This is in contrast to the stellar mass density, which grows substantially over that interval, highlighting the transient nature of the HI reservoir around galaxies. Today, the majority of the cosmic HI mass is in galaxies with $10^9 \lesssim M_{\text{HI}} \lesssim 10^{10} M_{\odot}$.

- We briefly examine numerical resolution convergence over a modest span of a factor of 8 in mass resolution between two simulations with momentum-driven winds. The HI richness shows better convergence than the HI mass function, which increases as the resolution is improved, equivalent to an increase in HI mass by a factor of 1.5. Hence resolution convergence has not been demonstrably achieved yet, but the differences remain comparable to or smaller than the statistical uncertainties.

Simultaneously matching the observed low mass end of the galactic stellar mass function and the low mass end slope of the HIMF has remained elusive in simulations and even in semi-analytic models (SAMs) of galaxy formation where one has much more freedom to adjust the physical model (Mo et al. 2005; Lu et al. 2012, 2013). Fundamentally, in most models this arises because the low mass end slope of the dark matter mass function is quite steep, which is exacerbated by low-mass galaxies being more HI-rich. Mo et al. (2005) argue that it is not possible to match the low mass end of the HIMF in a very broad range of CDM based galaxy formation models if one only makes two basic assumptions: First, that the original gas distribution has the same specific angular momentum as the dark matter halo, which it conserves as it forms an exponential disk (as in Mo, Mao, & White 1998); second, that stars only form in gas above a critical surface density Σ_{crit} , as observed (Kennicutt 1989). Almost all SAMs make these same assumptions. They populate some fraction of the cold dark matter halos with exponential gas disks and assume that any region originally above Σ_{crit} has its surface density lowered to Σ_{crit} . Making the (conservative) assumption that half of the gas is neutral, they overproduce the observed HIMF at the low mass end by a factor of more than five. Lu et al. (2012) confirm this result by adjusting all their SAM parameters using a Bayesian inference technique to match the observed low redshift K -band galaxy luminosity function, and find that they overpredict the observed HIMF by a large factor.

Lu et al. (2013) expand on this work, by adjusting their parameters to simultaneously match both the observed low- z K -band galaxy luminosity function and the HIMF. They find that this is not possible using standard models including the above assumptions, but it is possible if one does any of the following: 1) reduce Σ_{crit} by about an order of magnitude, 2) allow the disk to maintain an exponential profile while stars are forming, i.e. allow gas from large radii beyond the star forming radius to lose angular momentum by some process and move inwards, 3) include some pro-

cess that preheats the gas before it can be accreted into a dark matter halo or have the formula by which one determines the expected gas accretion rate be different for small mass halos. The first option violates observations by a substantial margin, and furthermore for smaller, lower-metallicity galaxies one observes Σ_{crit} to increase, not decrease (Bolatto, Wolfire, & Leroy 2013). Our simulations assume a three dimensional critical star formation density of $n_H = 0.13 \text{ cm}^{-3}$, which approximately matches the observed Σ_{crit} (Springel & Hernquist 2003), and hence this is not likely to be the difference between our models and the SAMs. The third possibility is also unlikely to be happening in our models, since we do not explicitly add any sort of preheating, nor do we explicitly vary the dark matter accretion rates from the usual Λ CDM expectations (e.g. Neistein & Dekel 2008; Faucher-Giguere et al. 2011). Hence if the arguments presented in Lu et al. (2013) apply to our simulations, then since we simultaneously match the HIMF and the stellar mass function, we suspect that angular momentum transport processes, either physical or numerical, cause the gas in our simulated disks to move inwards. We intend to investigate this and other possibilities in future work.

With the emergence of numerous major new radio facilities such as the Jansky Very Large Array (JVLA), MeerKAT (Karoo Array Telescope), and the Australian Square Kilometer Array Pathfinder (ASKAP), and eventually the SKA itself in the next decade, the future of H_I studies looks promising. Simulations like the ones presented here, hopefully improved substantially in the coming years, will be crucial for providing a context to these observations within the broader multi-wavelength landscape of galaxy evolution. At this point, it seems equally important both to probe the evolution of the H_I content of galaxies and to go deeper around nearby galaxies to better understand how H_I connects to the Cosmic Web. Such data promise to provide interesting new constraints on the key processes of galaxy evolution such as inflows, outflows, and wind recycling. This paper represents a first step towards providing a comprehensive framework for interpreting the wealth of forthcoming H_I data.

ACKNOWLEDGEMENTS

The authors acknowledge ... The simulations used here were run on the University of Arizona's SGI cluster, ice. This work was supported by the National Science Foundation under grant numbers AST-0847667, AST-0907998, AST-0908334, AST-0907651, and NASA grants NNX12AH86G and NNX10AJ95G. Computing resources were obtained through grant number DMS-0619881 from the National Science Foundation.

REFERENCES

Andrews, B. & Martini, P. 2013, ApJ, submitted, arXiv: Asplund, M., Grevesse, N., Sauval, A. J., Scott, P. 2009, ARA&A, 47, 481

Baldry, I. K., Glazebrook, K., Driver, S. P. 2008, MNRAS, 388, 945

Battisti, A. J. et al. 2012, ApJ, 744, 93

Becker, G.D., Hewett, P.C., Worseck, G., Prochaska, J.X. 2012, MNRAS, submitted, arXiv:1208.2584

Bell, E. F., Zheng, X. Z., Papovich, C., Borch, A., Wolf, C., Meisenheimer, K. 2007, ApJ, 663, 834

Blanton, M. R., Eisenstein, D., Hogg, D. W., Zehavi, I. 2006, ApJ, 645, 977

Blitz, L. & Rosolowsky, E. 2006, ApJ, 650, 933

Bolatto, A. D., Wolfire, M., Leroy, A. K. 2013, ARA&A, 51, in press, arXiv:1301.3498

Booth, C. M. & Schaye, J. 2009, MNRAS, 398, 53

Bouche, N. et al. 2010, ApJ, 718, 1001

Catinella, B. et al. 2010, MNRAS, 403, 683

Catinella, B. et al. 2012, A&A, 544, 65

Chabrier G., 2003, PASP, 115, 763

Cortese, L., Catinella, B., Boissier, S., Boselli, A., Heinis, S. 2011, MNRAS, 415, 1797

Daddi, E. et al. 2007, ApJ, 670, 156

Dalla Vecchia, C., Schaye, J. 2008, MNRAS, 387, 1431

Davé, R. 2008, MNRAS, 385, 147

Davé, R., ASPC, 419, 347

Davé, R., Oppenheimer, B. D., Katz, N., Kollmeier, J. A., Weinberg, D. H. 2010, MNRAS, 408, 2051

Davé, R., Oppenheimer, B. D., Finlator, K. M. 2011, MNRAS, 415, 11

Davé, R., Finlator, K. M., Oppenheimer, B. D. 2011, MNRAS, 416, 1354

Davé, R., Finlator, K. M., Oppenheimer, B. D. 2012, MNRAS, in press

Dekel, A., Birnboim, Y. 2006, MNRAS, 368, 2

Dekel, A. et al. 2009, Nature, 457, 451

Di Matteo, T., Springel, V., Hernquist, L. 2005, Nature, 433, 604

Duffy, A.R., Kay, S.T., Battye, R.A., Booth, C.M., Dalla Vecchia, C., Schaye, J. 2012, MNRAS, 420, 2799

Ellison, S. L., Patton, D. R., Simard, L., McConnachie, A. W. 2008, ApJL, 672, L107

Fardal, M. A., Katz, N., Weinberg, D. H., Davé, R., Katz, N. 2007, MNRAS, 379, 985

Faucher-Giguere, C. A., Kereš, D., Dijkstra, M., Hernquist, L., Zaldarriaga, M. 2010, ApJ, 725, 633

Faucher-Giguere, C. A., Kereš, D., Ma, C.-P. 2011, 417, 2982

Finlator, K., Davé, R., Papovich, C., Hernquist, L. 2006, ApJ, 639, 672

Finlator, K. & Davé, R. 2008, MNRAS, 385, 2181

Gabor, J. M., Davé, R., Oppenheimer, B. D., Finlator, K. M. 2011, MNRAS, 417, 2676

Gabor, J. M. & Davé, R. 2012, MNRAS, submitted

Geach, J. E., Smail, I., Moran, S. M., MacArthur, L. A., Lagos, C.d.P., Edge, A.C. 2011, ApJ, 730, 19

Giovanelli, R. 2005, AJ, 130, 2598

Haardt, F. & Madau, P. 2001, in proc. XXXVII Rencontres de Moriond, eds. D.M. Neumann & J.T.T. Van.

Haynes, M. P., Giovanelli, R., Chincarini, G. L. 1984, ARA&A, 22, 445

Haynes, M. P. et al. 2011, AJ, 142, 170

Hinshaw, G. et al. 2009, ApJS, 180, 225

Holwerda, B. W. et al. 2012, in proc. IAU Symposium 284, "The Spectral Energy Distribution of Galaxies" (SED2011), 5-9 Sep 2011, Preston, UK, eds. R.J. Tuffs & C.C. Popescu

Hopkins, A.M., Beacom, J.F. 2006, ApJ, 651, 142

Hopkins, P. F., Quataert, E., Murray, N. 2012, MNRAS, 412, 3522

Huang, S., Haynes, M. P., Giovanelli, R., Brinchmann, J., Stierwalt, S., Neff, S. G. 2012, AJ, 143, 133

Kassin, S.A. et al. 2007, ApJ, 660, L35

Katz, N., Weinberg, D. H., Hernquist, L. 1996, ApJS, 105, 19

Kennicutt, R. C. 1989, ApJ, 344, 685

Kennicutt, R. C. 1998, ApJ, 498, 541

Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, MNRAS, 363, 2

Kereš, D., Katz, N., Fardal, M., Davé, R., Weinberg, D. H. 2009, MNRAS, 395, 160

Kirkman, D., Tytler, D., Lubin, D., Charlton, J. 2007, MNRAS, 376, 1227

Lara-López, M. A. et al. 2010, A& 521, L53

Leitner, S. N. 2012, ApJ, 745, 149

Leroy, A. K., Walter, F., Brinks, E., Bigiel, F., de Blok, W. J. G., Madore, B., Thornley, M. D. 2008, AJ, 136, 2782

Lu, Y., Mo, H. J., Katz, N., Weinberg, M. D. 2012, MNRAS, 421, 1779

Lu, Y., Mo, H. J., Katz, N., Weinberg, M. D. 2012, MNRAS, submitted

Mannucci, F., Cresci, G., Maiolino, R., Marconi, A., Gnerucci, A. 2010, MNRAS, submitted, arXiv:1005.0006

McGaugh, S.S., Schombert, J.M., de Blok, W.J.G., Ziegler, M.J. 2010, ApJ, 708, L14

McGaugh, S.S. 2012, AJ, 143, 40

McKee, C. F. & Ostriker, J. P. 1977, ApJ, 218, 148

Meyer, M. J. et al. 2004, MNRAS, 350, 1195

Mo, H. J., Mao, S., White, S. D. M. 1998, MNRAS, 295, 319

Mo, H. J., Yang, X., van den Bosch, F.C., & Katz, N. 2005, MNRAS, 363, 1155

Moldar, S.M., Hearn, N., Haiman, Z., Bryan, G., Evrad, A.E., Lake, G. 2009, ApJ, 696, 1640

Moran, S.M. et al. 2012, ApJ, 745, 66

Murray, N., Quatert, E., Thompson, T. A. 2005, ApJ, 618, 569

Murray, N., Quatert, E., Thompson, T. A. 2010, ApJ, 709, 191

Narayanan, D., Bothwell, M., Davé, R. 2012, MNRAS, submitted

Neistein, E., Dekel, A. 2008, MNRAS, 388, 1792

Oppenheimer, B. D., Davé, R., Kereš, D., Katz, N., Kollmeier, J. A., Weinberg, D. H. 2010, MNRAS, 406, 2325

Oppenheimer, B. D. & Davé, R. 2006, MNRAS, 373, 1265

Oppenheimer, B. D. & Davé, R. 2008, MNRAS, 387, 577

Popping, A., Davé, R., Braun, R., Oppenheimer, B. D. 2009, A&A, 504, 15

Prochaska, J.X., Wolfe, A.M. 2009, ApJ, 696, 1543

Rao, S.M., Turnshek, D.A., Nestor, D.B. 2006, ApJ, 636, 610

Robertson, P., Shields, G. A., Blanc, G. A. 2012, ApJ, 748, 48

Rupke, D. S., Veilleux, S., & Sanders, D. B. 2005, ApJS, 160, 115

Salim, S. 2007, ApJS, 173, 267

Schmidt, M. 1959, ApJ, 129, 243

Springel, V. & Hernquist, L. 2003, MNRAS, 339, 289

Springel, V. & Hernquist, L. 2003, MNRAS, 339, 312

Springel, V. 2005, MNRAS, 364, 1105

Tacconi, L. J. et al. 2010, Nature, 463, 781

Tacconi, L. J. et al. 2012, ApJ, submitted, arXiv:1211.5743

Tremonti, C. A. et al. 2004, ApJ, 613, 898

Tully, R.B., Fisher, J.R. 1977, A&A, 54, 661

van de Voort, F., Schaye, J., Booth, C. M., Haas, M. R., Dalla Vecchia, C. 2010, MNRAS, submitted, arXiv:1011.2491

Verner, D. A. & Ferland, G. J. 1996, ApJS, 103, 467

Vogt, N.P., Haynes, M.P., Herter, T., Giovanelli, R. 2004, AJ, 127, 3273

Weinmann, S. M., Pasquali, A., Oppenheimer, B. D., Finlator, K., Mendel, J. T., Crain, R. A., Maccio, A. V. 2012, MNRAS, accepted

Wiersma, R.P., Schaye, K., Theuns, T., Dalla Vecchia, C., Tornatore, L. 2009, 399, 574

Zwaan, M. A., Meyer, M. J., Staveley-Smith, L., Webster, R. L. 2005, MNRAS, 359, L30